



Figure 6: Random samples from CIFAR-10 models. PixelCNN (Oord et al., 2016) and Flow++ samples reprinted from Ho et al. (2019), with permission.

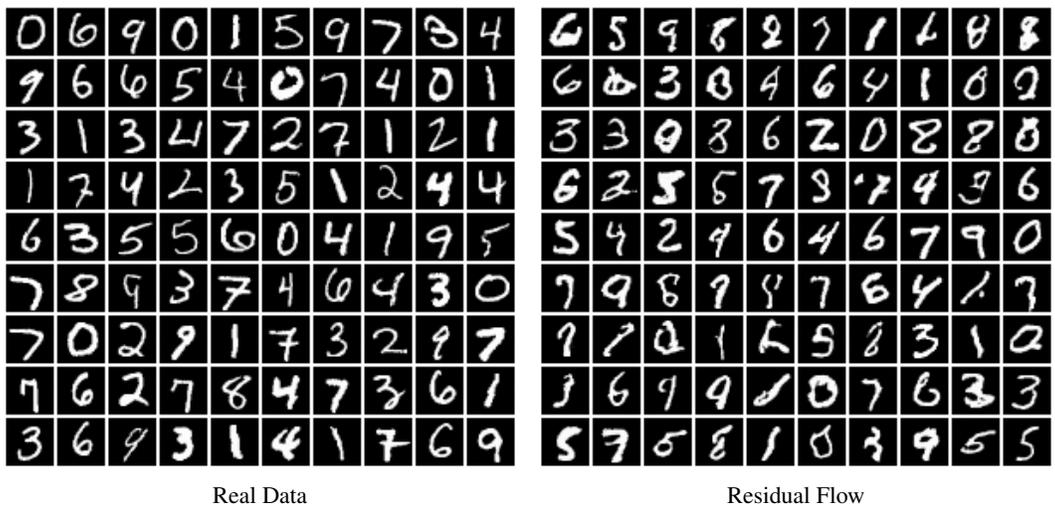


Figure 7: Random samples from MNIST.



Figure 8: Random samples from ImageNet 32x32.

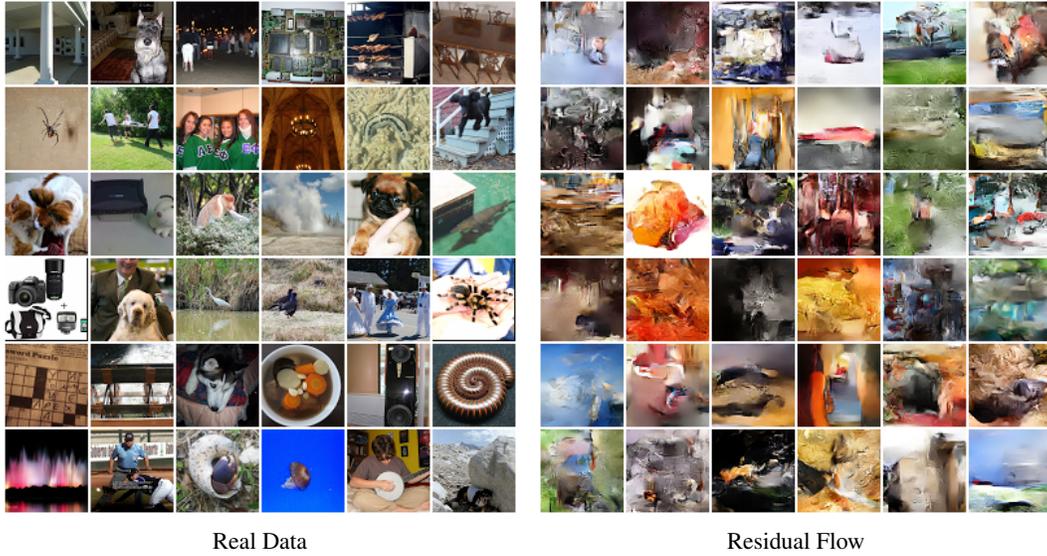


Figure 9: Random samples from ImageNet 64×64 .

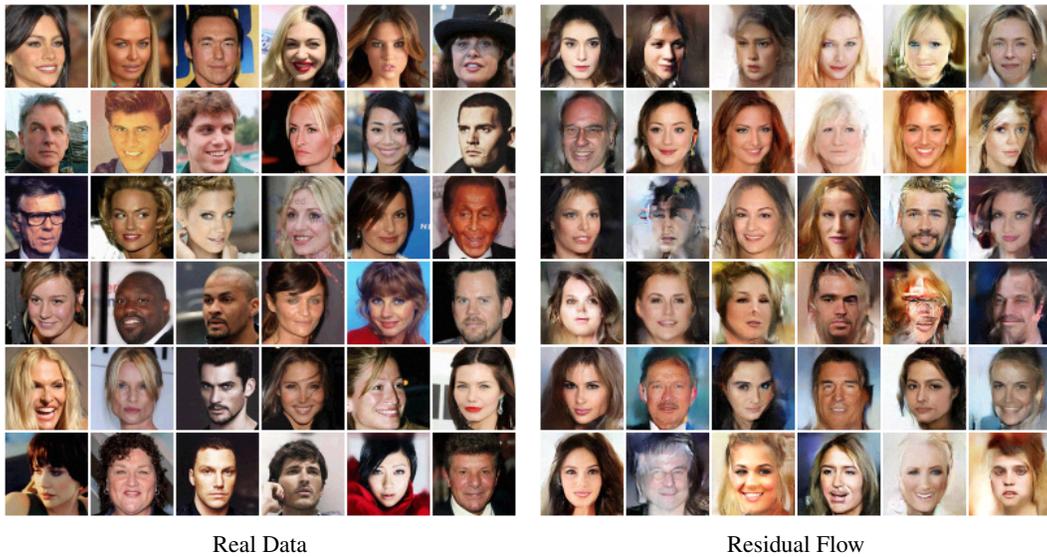


Figure 10: Random samples from 5bit CelebA 64×64 .

374 B Proofs

375 We start by formulating a Lemma, which gives the condition when the randomized truncated series is
 376 an unbiased estimator in a fairly general setting. Afterwards, we study our specific estimator and
 377 prove that the assumption of the Lemma is satisfied.

378 Note, that similar conditions have been stated in previous works, e.g. in [McLeish \(2011\)](#) and [Rhee
 379 and Glynn \(2012\)](#). However, we use the condition from ([Bouchard-Côté, 2018](#)), which only requires
 380 $p(N)$ to have sufficient support.

381 To make the derivations self-contained, we reformulate the conditions from ([Bouchard-Côté, 2018](#))
 382 in the following way:

383 **Lemma 3** (Unbiased randomized truncated series). *Let Y_k be a real random variable with*
 384 *$\lim_{k \rightarrow \infty} \mathbb{E}[Y_k] = a$ for some $a \in \mathbb{R}$. Further, let $\Delta_0 = Y_0$ and $\Delta_k = Y_k - Y_{k-1}$ for $k \geq 1$.*

385 *Assume*

$$\mathbb{E} \left[\sum_{k=0}^{\infty} |\Delta_k| \right] < \infty$$

386 *and let N be a random variable with support over the positive integers and $n \sim p(N)$. Then for*

$$Z = \sum_{k=0}^n \frac{\Delta_k}{\mathbb{P}(N \geq k)},$$

387 *it holds*

$$a = \lim_{k \rightarrow \infty} \mathbb{E}[Y_k] = \mathbb{E}_{n \sim p(N)}[Z] = a.$$

388 *Proof.* First, denote

$$Z_M = \sum_{k=0}^M \frac{\mathbb{1}[N \geq k] \Delta_k}{\mathbb{P}(N \geq k)} \quad \text{and} \quad B_M = \sum_{k=0}^M \frac{\mathbb{1}[N \geq k] |\Delta_k|}{\mathbb{P}(N \geq k)},$$

389 where $|Z_M| \leq B_M$ by the triangle inequality. Since B_M is non-decreasing, the monotone convergence theorem allows swapping the expectation and limit as $\mathbb{E}[B] = \mathbb{E}[\lim_{M \rightarrow \infty} B_M] =$
390 $\lim_{M \rightarrow \infty} \mathbb{E}[B_M]$. Furthermore, it is

$$\begin{aligned} \mathbb{E}[B] &= \lim_{M \rightarrow \infty} \mathbb{E}[B_M] = \lim_{M \rightarrow \infty} \sum_{k=0}^M \mathbb{E} \left[\frac{\mathbb{1}[N \geq k] |\Delta_k|}{\mathbb{P}(N \geq k)} \right] \\ &= \lim_{M \rightarrow \infty} \sum_{k=0}^M \frac{\mathbb{P}(N \geq k) \mathbb{E}|\Delta_k|}{\mathbb{P}(N \geq k)} = \mathbb{E} \left[\lim_{M \rightarrow \infty} \sum_{k=0}^M |\Delta_k| \right] < \infty, \end{aligned}$$

392 where the assumption is used in the last step. Using the above, the dominated convergence theorem
393 can be used to swap the limit and expectation for Z_M . Using similar derivations as above, it is

$$\mathbb{E}[Z] = \lim_{M \rightarrow \infty} \mathbb{E}[Z_M] = \lim_{M \rightarrow \infty} \mathbb{E} \left[\sum_{k=0}^M \Delta_k \right] = \lim_{M \rightarrow \infty} \mathbb{E}[Y_k] = a,$$

394 where we used the definition of Y_M and $\lim_{k \rightarrow \infty} \mathbb{E}[Y_k] = a$. □

395 *Proof. (Theorem 1)*

396 To simplify notation, we denote $J := J_g(x)$. Furthermore, let

$$Y_N = \mathbb{E}_v \left[\sum_{k=1}^N \frac{(-1)^{k+1}}{k} v^T J^k v \right]$$

397 denote the real random variable and let $\Delta_0 = Y_0$ and $\Delta_k = Y_k - Y_{k-1}$ for $k \geq 1$, as in Lemma 3. To
398 prove the claim of the theorem, we can use Lemma 3 and we only need to prove that the assumption
399 $\mathbb{E}_v[\sum_{k=1}^{\infty} |\Delta_k|] < \infty$ holds for this specific case.

400 In order to exchange summation and expectation via Fubini's theorem, we need to proof that
401 $\sum_{k=1}^{\infty} |\Delta_k| < \infty$ for all $v \sim \mathcal{N}(0, I)$. Using the assumption $\text{Lip}(g) < 1$, it is

$$\begin{aligned} \sum_{k=1}^{\infty} |\Delta_k| &= \sum_{k=1}^{\infty} \left| \frac{(-1)^{k+1}}{k} v^T J^k v \right| = \sum_{k=1}^{\infty} \frac{\|v^T J^k v\|_2}{k} \leq \sum_{k=1}^{\infty} \frac{\|v^T\|_2 \|J^k\|_2 \|v\|_2}{k} \\ &\leq 2\|v\|_2 \sum_{k=1}^{\infty} \frac{\|J\|_2^k}{k} \leq 2\|v\|_2 \sum_{k=1}^{\infty} \frac{\text{Lip}(g)_2^k}{k} = 2\|v\|_2 \log(1 - \text{Lip}(g)) < \infty, \end{aligned}$$

402 for an arbitrary v . Hence,

$$\mathbb{E}_v \left[\sum_{k=1}^{\infty} |\Delta_k| \right] = \sum_{k=1}^{\infty} \mathbb{E}_v[|\Delta_k|]. \tag{13}$$

403 Since $\text{tr}(A) = \mathbb{E}_v[v^T A v]$ for $v \sim \mathcal{N}(0, I)$ via the Skilling-Hutchinson trace estimator (Hutchinson,
404 1990; Skilling, 1989), it is

$$\mathbb{E}_v[|\Delta_k|] = \left| \frac{\text{tr}(J^k)}{k} \right|.$$

405 To show that (13) is bounded, we derive the bound

$$\frac{1}{k} |\text{tr}(J^k)| \leq \frac{1}{k} \left| \sum_{i=d}^d \lambda_i(J^k) \right| \leq \frac{1}{k} \sum_{i=d}^d |\lambda_i(J^k)| \leq \frac{d}{k} \rho(J^k) \leq \frac{d}{k} \|J^k\|_2 \leq \frac{d}{k} \text{Lip}(g)^k,$$

406 where $\lambda(J^k)$ denote the eigenvalues and $\rho(J^k)$ the spectral radius. Inserting this bound into (13) and
407 using $\text{Lip}(g) < 1$ yields

$$\mathbb{E}_v[|\Delta_k|] \leq d \sum_{k=1}^{\infty} \frac{\text{Lip}(g)^k}{k} = -d \log(1 - \text{Lip}(g)) < \infty.$$

408 Hence, the assumption of Lemma 3 is verified. \square

409 *Proof. (Theorem 2)*

410 The result can be proven in an analogous fashion to the proof of Theorem 1, which is why we only
411 present a short version without all steps.

412 By obtaining the bound

$$\begin{aligned} \sum_{k=0}^{\infty} \left| (-1)^k v^T \left(J(x, \theta)^k \frac{\partial(J_g(x, \theta))}{\partial \theta} \right) v \right| &\leq 2 \|v\|_2 \left\| \frac{\partial(J_g(x, \theta))}{\partial \theta} \right\| \sum_{k=0}^{\infty} \text{Lip}(g)^k \\ &= 2 \|v\|_2 \left\| \frac{\partial(J_g(x, \theta))}{\partial \theta} \right\| \frac{1}{1 - \text{Lip}(g)} < \infty, \end{aligned}$$

413 Fubini's theorem can be applied to swap the expectation and summation. Furthermore, by using the
414 trace estimation and similar bounds as in the proof of Theorem 1, the assumption $\mathbb{E}[\sum_{k=0}^{\infty} |\Delta_k|] < \infty$
415 from Lemma 3 can be proven. \square

416

417 C Memory-Efficient Gradient Estimation of Log-Determinant

418 Derivation of gradient estimator via differentiating power series:

$$\begin{aligned} \frac{\partial}{\partial \theta_i} \log \det(I + J_g(x, \theta)) &= \frac{\partial}{\partial \theta_i} \left(\sum_{k=1}^{\infty} (-1)^{k+1} \frac{\text{tr}(J_g(x, \theta)^k)}{k} \right) \\ &= \text{tr} \left(\sum_{k=1}^{\infty} \frac{(-1)^{k+1}}{k} \frac{\partial(J_g(x, \theta)^k)}{\partial \theta_i} \right) \end{aligned}$$

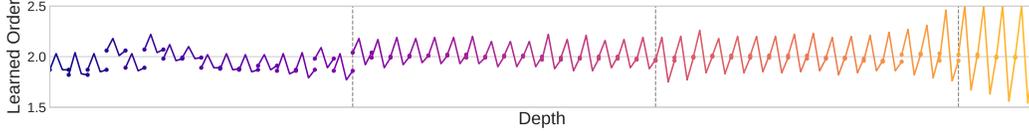


Figure 11: **Learned norm orders on CIFAR-10.** Each residual block is visualized as a single line. The input and two hidden states for each block use different normed spaces. We observe multiple trends: (i) the norms for the first hidden states are consistently higher than the input, and lower for the second. (ii) The orders for the hidden states drift farther away from 2 as depth increases. (iii) The ending order of one block and the starting order of the next are generally consistent and close to 2.

419 Derivation of memory-efficient gradient estimator:

$$\begin{aligned} & \frac{\partial}{\partial \theta_i} \log \det (I + J_g(x, \theta)) \\ &= \frac{1}{\det(I + J_g(x, \theta))} \left[\frac{\partial}{\partial \theta_i} \det (I + J_g(x, \theta)) \right] \end{aligned} \quad (14)$$

$$= \frac{1}{\det(I + J_g(x, \theta))} \left[\det(I + J_g(x, \theta)) \operatorname{tr} \left((I + J(x, \theta))^{-1} \frac{\partial (I + J_g(x, \theta))}{\partial \theta_i} \right) \right] \quad (15)$$

$$\begin{aligned} &= \operatorname{tr} \left((I + J(x, \theta))^{-1} \frac{\partial (I + J_g(x, \theta))}{\partial \theta_i} \right) \\ &= \operatorname{tr} \left((I + J(x, \theta))^{-1} \frac{\partial (J_g(x, \theta))}{\partial \theta_i} \right) \\ &= \operatorname{tr} \left(\left[\sum_{k=0}^{\infty} (-1)^k J(x, \theta)^k \right] \frac{\partial (J_g(x, \theta))}{\partial \theta_i} \right) \end{aligned} \quad (16)$$

420 Note, that (14) follows from the chain rule of differentiation, for the derivative of the determinant in
421 (15), see (Petersen and Pedersen, 2012) (eq. 46). Furthermore, (16) follows from the properties of a
422 Neumann-Series which converges due to $\|J_g(x, \theta)\| < 1$.

423 Hence, if we are able to compute the trace exactly, both approaches will return the same values for
424 a given truncation n . However, when estimating the trace via the Hutchinson trace estimator the
425 estimation is not equal in general:

$$v^T \left(\sum_{k=1}^{\infty} \frac{(-1)^{k+1}}{k} \frac{\partial (J_g(x, \theta)^k)}{\partial \theta_i} \right) v \neq v^T \left(\left[\sum_{k=0}^{\infty} (-1)^k J_g^k(x, \theta) \right] \frac{\partial (J_g(x, \theta))}{\partial \theta_i} \right) v.$$

426 Another difference between both approaches is their memory consumption of the corresponding
427 computational graph. The summation $\sum_{k=0}^{\infty} (-1)^k J_g^k(x, \theta)$ is not being tracked for the gradient,
428 which allows to compute the gradient with constant memory (constant with respect to the truncation
429 n).

430 D Generalized Spectral Norm

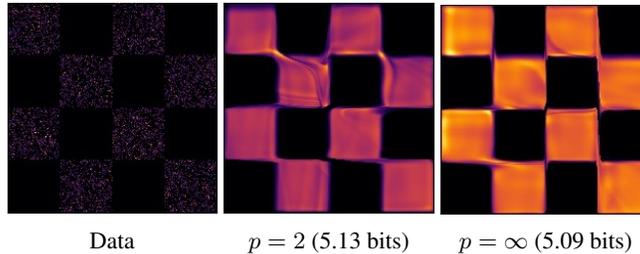


Figure 12: Learned densities on Checkerboard 2D.

431 **Using different induced p -norms on Checkerboard 2D.** We experimented with the checkerboard
 432 2D dataset, which is a rather difficult two-dimensional data to fit a flow-based model on due to the
 433 discontinuous nature of the true distribution. We used brute-force computation of the log-determinant
 434 for change of variables (which, in the 2D case, is faster than the unbiased estimator). In the 2D case,
 435 we found that ∞ -norm always outperforms or at least matches the $p = 2$ norm (ie. spectral norm).
 436 Figure 12 shows the learned densities with 200 residual blocks. The color represents the magnitude of
 437 $p_\theta(x)$, with brighter values indicating larger values. The ∞ -norm model produces density estimates
 438 that are more evenly spread out across the space, whereas the spectral norm model focused its density
 439 to model between-density regions.

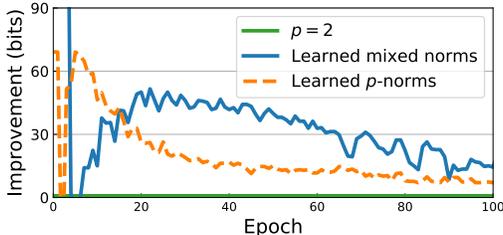


Figure 13: Improvement from using generalized spectral norm on CIFAR-10.

440 **Learning norm orders on CIFAR-10.** We used $1 + \tanh(s)/2$ where s is a learned weight. This
 441 bounds the norm orders to $(1.5, 2.5)$. We tried two different setups. One where all norm orders are
 442 free to change (conditioned on them satisfying the constraints (11)), and another setting where each
 443 states within each residual block share the same order. Figure 13 shows the improvement in bits from
 444 using learned norms. The gain in performance is marginal, and the final models only outperformed
 445 spectral norm by around 0.003 bits/dim. Interestingly, we found that the learned norms stayed around
 446 $p = 2$, shown in Figure 11, especially for the input and output spaces of g , ie. between blocks. This
 447 may suggest that spectral norm, or a norm with $p = 2$ is already optimal in this setting.

448 E Experiment Setup

449 We use the standard setup of passing the data through a “unsquashing” layer (we used the logit
 450 transform (Dinh et al., 2017)), followed by alternating multiple blocks and squeeze layers (Dinh et al.,
 451 2017). We use activation normalization (Kingma and Dhariwal, 2018) before and after every residual
 452 block. Each residual connection consists of

$$453 \quad \text{LipSwish} \rightarrow 3 \times 3 \text{ Conv} \rightarrow \text{LipSwish} \rightarrow 1 \times 1 \text{ Conv} \rightarrow \text{LipSwish} \rightarrow 3 \times 3 \text{ Conv}$$

454 with hidden dimensions of 512. Below are the architectures for each dataset.

455 **MNIST.** With $\alpha = 1e-5$.

$$456 \quad \text{Image} \rightarrow \text{LogitTransform}(\alpha) \rightarrow 16 \times \text{ResBlock} \rightarrow [\text{Squeeze} \rightarrow 16 \times \text{ResBlock}] \times 2$$

457 **CIFAR-10.** With $\alpha = 0.05$.

$$458 \quad \text{Image} \rightarrow \text{LogitTransform}(\alpha) \rightarrow 16 \times \text{ResBlock} \rightarrow [\text{Squeeze} \rightarrow 16 \times \text{ResBlock}] \times 2$$

459 **ImageNet 32x32.** With $\alpha = 0.05$.

$$460 \quad \text{Image} \rightarrow \text{LogitTransform}(\alpha) \rightarrow 32 \times \text{ResBlock} \rightarrow [\text{Squeeze} \rightarrow 32 \times \text{ResBlock}] \times 2$$

461 **ImageNet 64x64.** With $\alpha = 0.05$.

$$462 \quad \text{Image} \rightarrow \text{Squeeze} \rightarrow \text{LogitTransform}(\alpha) \rightarrow 32 \times \text{ResBlock} \rightarrow [\text{Squeeze} \rightarrow 32 \times \text{ResBlock}] \times 2$$

463 **CelebA 64×64.** With $\alpha = 0.05$.

464 Image \rightarrow Squeeze \rightarrow LogitTransform(α) \rightarrow 16×ResBlock \rightarrow [Squeeze \rightarrow 16×ResBlock] ×3

465 For hybrid modeling on CIFAR-10, we replaced the logit transform with normalization by the
466 standard preprocessing of subtracting the mean and dividing by the standard deviation across the
467 training data. The MNIST and SVHN architectures for hybrid modeling were the same as those for
468 density modeling.

469 For augmenting our flow-based model with a classifier in the hybrid modeling experiments, we added
470 an additional branch after every squeeze layer and at the end of the network. Each branch consisted
471 of

472 3×3 Conv \rightarrow ActNorm \rightarrow ReLU \rightarrow AdaptiveAveragePooling((1, 1))

473 where the adaptive average pooling averages across all spatial dimensions and resulted in a vector of
474 dimension 256. The outputs at every scale were concatenated together and fed into a linear softmax
475 classifier.

476 **Adaptive number of power iterations.** To account for variable weight updates during training,
477 we used an adaptive version of spectral normalization for convolutional layers (Gouk et al., 2018)
478 where we performed as many iterations as needed until the relative change in the estimated spectral
479 norm was sufficiently small. As this also reduced the number of iterations when no weight changes
480 occur, it resulted in speed comparable to always performing 5 iterations of power method.

481 **Optimization.** For stochastic gradient descent, we used Adam (Kingma and Ba, 2014) with a
482 learning rate of 0.001 with otherwise default hyperparameters. We used Polyak averaging (Polyak
483 and Juditsky, 1992) for evaluation with a momentum of 0.999.

484 **Preprocessing.** For density estimation experiments, we used random horizontal flipping for CIFAR-
485 10 and CelebA.

486 For hybrid modeling and classification experiments, we used random cropping after reflection padding
487 with 4 pixels for SVHN and CIFAR-10. CIFAR-10 also included random horizontal flipping.